REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggesstions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any oenalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
	Technical Report	-

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
NEMS (Nanoelectromechanicsl Systems) Networks:A Novel Validation Platform for Controlling Interconnected Dynamical Networks	W911NF-13-1-0340
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
	611103

6. AUTHORS	5d. PROJECT NUMBER
Matthew Matheny, Michael Roukes, James Crutchfield, Raissa D'Souza, Warren Fon	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES	8. PERFORMING ORGANIZATION REPORT NUMBER
University of California - Davis 1850 Research Park Drive Suite 300 Davis, CA 95618 -6153	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
	63826-NS-MUR.48

12. DISTRIBUTION AVAILIBILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

We review the experimental role NanoElectroMechanical System (NEMS) networks play vis a vis the MURI's scientific and validation goals which target new theoretical understanding of the dynamics of large-scale interconnected networks and their control. Our plans for NEMS networks are substantially updated based on our theoretical and experimental results during the MURI's first two years, providing a clear roadmap for the coming years that discusses relative benefits and costs. For example, we briefly compare how NEMS networks compare to alternative, potential validation platforms, including one based on Field Programmable Gate Arrays (FPGAs). We

15. SUBJECT TERMS
Nanoelectromechanical Systems, synchronization, control

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Raissa D'Souza
UU	UU	UU	UU		19b. TELEPHONE NUMBER
					530-754-8405

Standard Form 298 (Rev 8/98)
Prescribed by ANSI Std. Z39.18

Report Title

NEMS (Nanoelectromechanicsl Systems) Networks:A Novel Validation Platform for Controlling Interconnected Dynamical Networks

ABSTRACT

We review the experimental role NanoElectroMechanical System (NEMS) networks play vis a vis the MURI's scientific and validation goals which target new theoretical understanding of the dynamics of large-scale interconnected networks and their control. Our plans for NEMS networks are substantially updated based on our theoretical and experimental results during the MURI's first two years, providing a clear roadmap for the coming years that discusses relative benefits and costs. For example, we briefly compare how NEMS networks compare to alternative, potential validation platforms, including one based on Field-Programmable Gate Arrays (FPGAs). We also provide a relatively detailed outline of system-on-chip implementation of large-scale NEMS networks that will allow us to explore the nonlinear dynamics, statistical mechanics, and control of very large oscillator networks.

NEMS (Nanoelectromechanicsl Systems) Networks:
A Novel Validation Platform for Controlling Interconnected Dynamical Networks

Warren Fon, Matthew Matheny, & Michael Roukes (Caltech)
James Crutchfield & Raissa D'Souza (UCD)

Supplement to the August 2015
MURI Interim Progress Report

Predicting and Controlling Systems of Interdependent Networks:
Exploiting interdependence for control

Technical POC: Professor Raissa D'Souza
Depts. of Computer Science & of Mechanical and Aerospace Engineering
Complexity Sciences Center
University of California at Davis
1 Shields Avenue, Davis CA 95616

Telephone: 530-220-3069 Fax: 530-752-4767
Email: raissa@cse.ucdavis.edu

Principal Investigators:
Raissa D'Souza, UC Davis, Lead
James P. Crutchfield, UC Davis
Leonardo Dueñas-Osorio, Rice University
Mehran Mesbahi, University of Washington
Michael Roukes, California Institute of Technology

NEMS (Nanoelectromechanicsl Systems) Networks:
A Novel Validation Platform for Controlling Interconnected Dynamical Networks

Warren Fon, Matthew Matheny, & Michael Roukes (Caltech)
James Crutchfield & Raissa D'Souza (UCD)

Abstract

We review the experimental role NanoElectroMechanical System (NEMS) networks play vis a vis the MURI's scientific and validation goals which target new theoretical understanding of the dynamics of large-scale interconnected networks and their control. Our plans for NEMS networks are substantially updated based on our theoretical and experimental results during the MURI's first two years, providing a clear roadmap for the coming years that discusses relative benefits and costs. For example, we briefly compare how NEMS networks compare to alternative, potential validation platforms, including one based on Field-Programmable Gate Arrays (FPGAs). We also provide a relatively detailed outline of system-on-chip implementation of large-scale NEMS networks that will allow us to explore the nonlinear dynamics, statistical mechanics, and control of very large oscillator networks.

Table of Contents

I. Why NEMS Networks?

Overview

The highly interdisciplinary approach of our NetControl MURI is centered on the development of analytical and numerical theories that are both informed and validated by empirical studies of real-world systems. Our focus is to develop new approaches that exploit network interdependence for network control. To develop robust theoretical principles and methods that can become widely applicable, it is essential to carry out these studies with well-characterized experimental test beds that can provide detailed information – down to the level of the dynamics of each individual node, and down to the "quantal" level of information exchange at each of the network's edges. Nanoscale mechanical oscillators, that is, nanoelectromechanical systems (NEMS), are an especially auspicious platform for these pursuits – as will be described below.

Our focus is on state-dependent networks, where states of both the system and the environment play critical roles. Accordingly, it is essential to develop an experimental platform for these efforts that can be upscaled to comprise a sufficiently large number of degrees-of-freedom that collective phenomena and emergent system properties are manifested. Further, real-world systems are always subject both to internal fluctuations, and to noisy and dynamic environments; hence, any meaningful experimental platform must be capable of including such aspects. Also essential is to be able to investigate the evolution of network dynamics under deterministic perturbations, complex signal injection, and detailed nodal and edge control. Rich topologies involving layered networks, networks of networks, and adaptive network topologies (a new domain for Network Science) provide especially fertile ground for these explorations.

As network size increases and the topology of network connectivity becomes complex, the temporal evolution of the system tends to span an increasingly vast hierarchy of time scales – with dynamics that often evolves over many decades in time. This poses a significant challenge: to elucidate the richness of the phenomena that are manifested we must track the detailed time evolution of the system – from its fastest "elemental" dynamics at the nodes, to the anomalous local-to-global diffusion of coherent and synchronized patterns that emerge. For systems with a large number of degrees-of-freedom it is notoriously difficult to capture glassy dynamics; these generally evolve very slowly, yet are completely dependent on the fast nodal dynamics for their emergence. Further, exploration of some of the most interesting and important classes of emergent behaviors induced by system heterogeneity can be essentially impossible to predict a priori. In this domain, experimental investigations – guided and then analyzed within theoretical frameworks – are key.

These general comments mask a very practical question: At what network size does one move from the low-dimensional behavioral regime in which tools from dynamical system theory for analyzing the geometry of state-space organization are appropriate to an approximate "thermodynamic limit" in which appeals to microscopic randomness allow statistical mechanics (its averaging) methods to come to the fore. In truth, this is a quite important contemporary theoretical problem and one to which we hope to contribute with our NEMS networks. However,

there is an immediate and compelling planning issue: How large should these networks be to allow us to study dynamical networks in the two regimes and to observe the transition between the regimes? Keeping the open nature of the question in mind, we would venture that 50 to 100 node networks will still be in the dynamical systems regime; whereas 1000-2000 node networks will effectively be in the thermodynamic regime. These ballpark figures inform the following discussion, in which we compare implementation platforms and lay out plans for developing NEMS networks in VLSI. Ballpark estimates aside, when considering what is an appropriate network size one can also appeal to the many real world systems with 1000 or more nodes – from distributed-power power grids to the brain – to which we hope our ultimate results and methods would apply.

In this light, we believe NEMS networks provide an unprecedented vehicle to achieve realistic, digital-artifact-free "analog simulations" of such phenomena. NEMS are intrinsically very-low-dissipation (high quality factor) resonators, and with them high-performance, high-frequency oscillator "nodes" can be constructed. Starting from this high performance level of the individual nodes, real physical dissipation and noise can be introduced – allowing the nodal dynamics to easily span regimes from weak to strong dissipation, and from coherent to strongly fluctuating dynamics. Access to virtually unlimited network topologies is possible; the varieties span from simple networks to the aforementioned networks-of-networks, layered topologies, and adaptive network topologies (which, for example, can be implemented via "smart" edges with local rules.)

There are several unique aspects of NEMS networks that are essential for this work. First, NEMS networks provide the capability of reliably capturing the full complexity of nonlinear dynamics. By this we mean that it is possible to impose detailed perturbations from local to global scales, while simultaneously observing the system response down to the unit level of dynamics at each nodes. Since "complete" observability of the system's internal dynamics is possible – over the aforementioned hierarchy of time scales – it becomes possible to search for overarching system observables, that is, order parameters, that best characterize the emergent phenomena such systems manifest. Second, NEMS are truly analog and continuous-time real-world systems. Modeling and capturing the true behavior of real-world systems has often proven to be elusive (even impossible!) through pursuits based on discrete numerical modeling or with digital electronic systems, given the sensitivity of complex system dynamics to minute perturbations.

Essential NEMS attributes for this work

NEMS resonators can span the frequency range from the kHz-scale to many GHz. This makes it straightforward to track fast system dynamics while observing the emergence of slow, even glassy, collective dynamics that evolve, for example, over many billions of the fast nodal cycles – and to achieve this in measurements intervals carried out over feasible laboratory (graduate researcher!) time scales. Further, NEMS are devices with both an ultralow intrinsic dissipation (high Q) and a low threshold onset of nonlinear dynamics. Q-factors in the range of 10^4 to $>10^6$ are possible, even at room temperature, whereas analog electrical circuits in this frequency range typically have Q's ~100 or less. With very modest drive levels above the intrinsic thermodynamic fluctuations of the nodes, strong nonlinearity can be induced. Together, these attributes make it

possible to completely control, moreover to tune, nodal parameters such as frequency, intrinsic nonlinearity, dissipation over a very wide range. Further, NEMS resonators are straightforward to configure into high-performance oscillator nodes that are extremely easy to couple controllably and deterministically. With our recent advanced in NEMS based on Silicon very-large-scale-integration (VLSI) it is now feasible to co-integrate them with CMOS transistors to have nodal gain necessary to create local feedback oscillators. With these "active" nodes extremely large network ensembles can be realized, with complete freedom to design arbitrary array architectures. Finally, by reducing ambient temperatures into the cryogenic realm, NEMS dynamics will transition from the classical to the quantum domain – opening as yet unexplored frontiers in quantum network science.

Our NEMS Networks roadmap from 30,000 feet

First phase (Y1-3)

The culmination of our NEMS efforts, which we expect to achieve in the second phase of our MURI effort, will involve development and exploration of coupled NEMS oscillators in networks upscaled to sizes involving a thousand, and possibly many thousands of, oscillators. Systems at this scale will give us access to the physics of disorder, to anomalous diffusion of order parameters and coherent patterns, to ultraslow glassy dynamics – phenomena that are only manifested in large networks. Further, with this number of nodes it will be possible to meaningfully assemble networks-of-networks and layered networks where the "constituent" networks dynamics are themselves large enough to manifest complex, even emergent, properties.

As detailed below, these experimental pursuits require Silicon VLSI, both for the NEMS and for the embedding CMOS circuitry – and this can only be carried out within the context of partnerships with $B-scale microelectronics foundries. **As this upscaling is both complex, costly, and involves design/fabrication cycle times that are somewhat long – the NEMS network efforts in the first phase of our MURI phase focus on printed-circuit-board NEMS networks, comprising <100 NEMS nodes, which can be implemented faster, as well as more easily and economically. These foundational first efforts are providing the initial guidance to allow us to pursue the larger NEMS networks (100 nodes and beyond) with confidence.**

Second phase (Y3-5)

The second phase of our explorations will focus on the aforementioned upscale networks with, potentially, thousands of NEMS oscillator nodes, with an even greater number edges connecting them. These complex nonlinear nodes and edges will provide a richness of phenomena manifested as limit cycles and basins of attraction that will be fertile ground for the development of novel network control strategies. Further, we intend to explore implementation of dynamical, *i.e.* "smart", edges through implementation of local rules programmed into the local CMOS-derived edge connections. This will open possibilities for exploring dynamical network topologies and, further, real-time strategies for adaptive topological control of networks. We believe this should be a new and compelling scientific frontier.

Culmination of our efforts with NEMS Networks: Next-generation Analog Simulation

We believe that our NEMS instantiations will provide a highly-configurable platform for complex analog simulations of layered networks. As the evolution of these systems' fast dynamics will be at high- to very-high-frequencies, their slow and emergent properties will evolve over reasonable (practical) laboratory timescales. As mentioned, we envision achieving complex networks involving a thousand nodes (or more), with dynamical elements and architectures that are completely configurable in real time. In reaching this level we will have moved from studying the nonlinear physics of one or several coupled NEMS oscillators to a mode in which they are seen as computing elements suitable for simulation of other systems of interest. And, then, once we have scaled up to networks consisting of 1000s of nodes, we will have in effect a powerful simulator for large-scale real world networks. We can imagine that practitioners would investigate their own application problems by programming (configuring) these NEMS network simulators.

II. Synergies and Differences with Alternate Validation Platforms

Although oscillator networks have been successfully implemented in other platforms, we believe that the NEMS instantiation provides compelling and transformational advantages. As we have previously mentioned, implementations based on analog electronics suffer from the inability of achieve ultralow dissipation and tunable nonlinear elements. These limitations, in fact, have motivated a current trend in frequency control industry toward use of microelectromechanical systems (MEMS)–based resonators as the basis for for electronic oscillators and clocks. Josephson (superconducting) junction arrays have been explored, but these must always be operated at cryogenic temperatures (below T_C, typ. few K), and their intrinsic fast dynamics are, in fact, too fast to be convenient and economically practical. Their dynamics spans the multi-GHz regime, the fundamental Josephson frequency voltage relation is h/2e = 484 MHz/µV. Voltage thresholds to achieve robust operation in the presence of thermal fluctuations generally must exceed V_{thresh} > 10's of µV. As the waveforms engendered are rich in harmonics (and preservation of a multiplicity of these harmonics is critical to the elemental behavior), ultra wideband measurement approaches and circuitry becomes essential.) Further, the intrinsic µV-scale operation of superconducting devices and their requisite low-temperature operation is generally a complete mismatch for conventional electronics. Finally, Josephson elements solely permit realization of rotators (phase oscillators)–and thereby restricts the richness of real-world phenomena that can be modeled.

Comparison of NEMS oscillators and FPGA Boolean Phase Oscillators (BPOs)

Recently, the Gauthier group has published two papers exploring synchronization of ring oscillators implemented in FPGAs. A ring oscillator's essential element is the inverter/buffer, which serves to add gain and delay, a sufficient condition for oscillation. The frequency of oscillation is where is the total delay through the inverter stages. This type of oscillator does not have 'amplitude'. It is what is known as a "phase oscillator" or a "rotator", just like a Josephson

Junction. The coupling is induced by making a phase comparison using an XOR gate (output 0 for the same value, output 1 for different values) and using the error signal to adjust the frequency.

'Small' Networks with Adjustable Coupling

In their paper Synchronization of Coupled Boolean Phase Oscillators PRE 2014, a single oscillator is synchronized to an external signal and two oscillators are coupled together. The oscillator (called here Boolean Phase Oscillator, BPO) implemented is actually a combination of two ring oscillators (with different frequencies) with a Boolean switch to rapidly switch between the two (see Figure 1). This rapid switching between the two will give an average frequency between the two frequencies. The phase difference between a BPO and an incoming signal produces the error bit (Vc in Figure 1b) which is sent to the Boolean switch of the two ring oscillators. This will shift average frequency of the BPO. The amount of shift in frequency for the BPO for the error signal is proportional to the frequency difference between the internal ring oscillators. The strength of the coupling is adjusted via the internal frequency difference between the two ring oscillators that make up a single BPO.
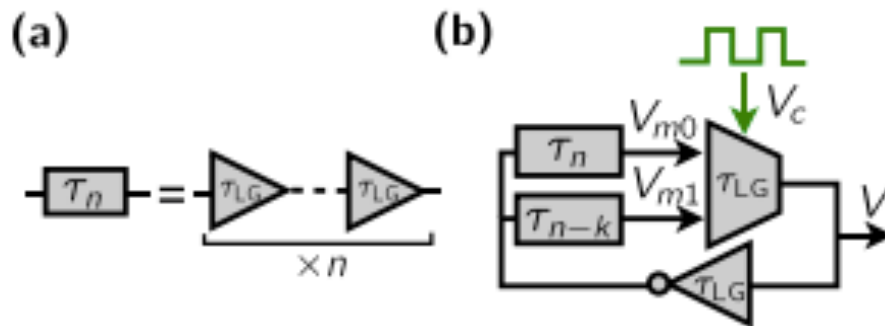

Figure 1. Boolean Phase Oscillator (BPO) Elements. a) Time delay from buffers b) BPO composed of two different ring oscillators.

'Large' Networks

In their paper, Transient Scaling and Resurgence of Chimera States in Coupled Boolean Phase Oscillators, PRE 2014 the BPO is implemented with one ring oscillator composed of 2R+1 delay stages, with each delay stage controlled by a Boolean switch (see Figure 2). The total frequency is determined by how many delay stages are turned on. Then as many as 128 BPOs are coupled together in a ring. For a single BPO coupled to 2R other BPOs, the delay stages are controlled by the error bit of the XOR phase comparator. Thus, if the phases between the BPO and the coupled BPOs are different, there is a digital shift in frequency. This is sufficient to demonstrate Kuramoto dynamics in the ring. There is no adjustment of coupling strength done in this paper.
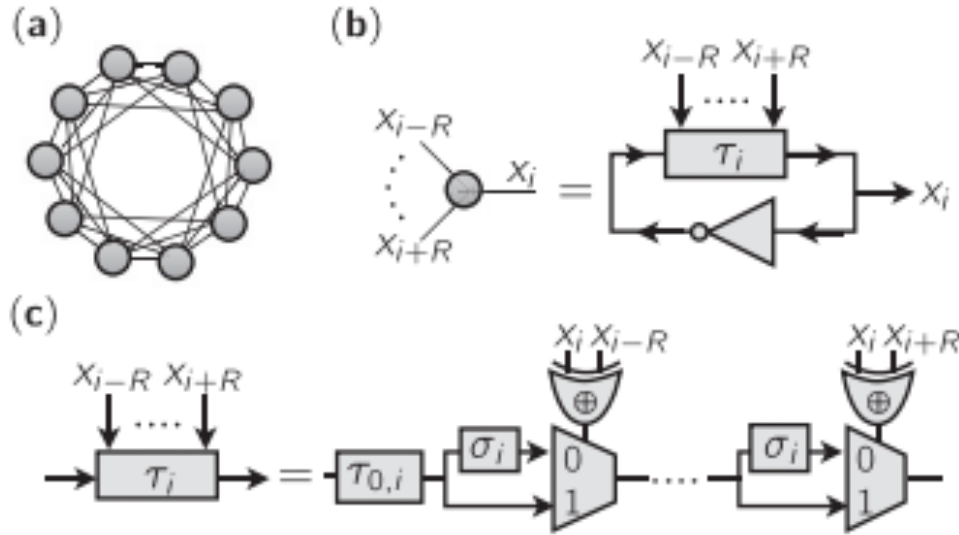
Figure 2. BPO network topologies. a) 'Ring' Network. b) Ring oscillator acting as BPO. c) Coupling works as the Boolean inputs to delay stages.

FPGA Scaling

The FPGA BPO platform is excellent for an inexpensive large scale demonstration of basic Kuramoto dynamics, but its not clear it can go beyond this limited space of synchronization. In addition, the adjustment of coupling and frequency (there is no nonlinearity here since there is no amplitude) must be done digitally since the bias voltage of the inverter stages cannot be adjusted. In order to change a parameter statically the configuration of the FPGA must be changed, which takes ~minutes. For dynamic network changes, a controllable n-bit stage must be set per BPO. For example, to set the frequency of a ring oscillator to 8-bit resolution, the delay path must have 256 delay stages each with a switch control. This means that for n bit resolution of a control, $2n^2$ elements have to be programmed. In addition, internal logic elements have to be set up to control this number of switches. To change the coupling, the same argument applies. The current BPOs in the 128 oscillator case required 218 logic elements each. Assuming that 8-bit resolution is desired on coupling and frequency, this requires ~1k logic elements/BPO. The maximum number of logic elements (LEs) for their current FPGA is ~100k. Currently the highest end FPGA ~1-5Million LEs. The cost of a 1 Millions element FPGA ~10k USD. This would set the maximum number of BPOs with dynamic coupling at ~1k.

NEMS oscillators

The NEMS system is different in that it is completely analog. The oscillators are complex-valued vectors in phase space. The behavior is not limited to only Kuramoto type synchronization, but is able to demonstrate true Ginzberg-Laudau diffusion on the network. Our PCB NEMS oscillator system is very flexible; each element can be modified. When moving to VLSI, the NEMS oscillator

system can support thousands of oscillators, in line with the calculations for ring oscillators implemented the best FPGAs currently on the market. Also, while there exists variation in the FPGA BPOs, there will exist both variation and noise in the NEMS. In a DAC, n-bit scaling is created by turning on resistors scaled by factors of 2, so n-bit DACs need ~2n not 2^{n^2} elements, reducing the number of elements necessary for precise control.

In short, the inherent analog nature of NEMS oscillators and so networks build from them is likely the mathematically and scientifically most distinctive difference with the inherently discrete-state, discrete time FPGAs and their networks. That said, we believe both provide excellent bases for probing the dynamics of large-scale networks. However, one must keep in mind that ultimately their dynamical phenomena will simply be different, even if there are similarities in restricted cases. Based on long experience with both discrete and continuous dynamical systems, likely the NEMS network being analog will generically exhibit richer phenomena.

Finally, that same experience suggests that we address a common question that arises whenever one is developing a new, hardware-specific platform for exploring dynamical behavior: Why not just simulate? And, of course, the question implicitly means simulation on conventional, perhaps cluster-based, digital computers. The first response is that we are already developing simulators for individual and coupled NEMS, including anticipating conventional simulators for NEMS networks. The second response is that, in a very real sense, the speed of FPGAs is an example of the fastest all-digital implementation, largely since they can be closely tailored to a given problem, whereas conventional digital computers must devote substantial chip resources to service and support general-purpose computation. To summarize, contemporary Linux clusters will be orders of magnitude slower and to get into comparable simulation-speed regimes and cost 1M$s. That is, clusters will not scale well to very large oscillator networks, especially if one's goals are to study analog oscillator networks.

Table I. Comparison between NEMS and FPGA Networks

Technology	NEMS-CMOS	FPGA
Node Dynamics	Continuous-time, Continuous state	Discrete-time, Discrete-state
Oscillator Type	NEMS + Electronic amplifier (analog)	Ring Oscillator (digital)
Representation	Amplitude + Phase	Phase
Frequency Band	HF/VHF	HF/VHF
Control Elements	Switched resistor networks	Switched delay buffers
Control Parameters	Coupling, Nonlinearity, Frequency	Frequency, Coupling
Network Topologies	Variable, on the fly	Variable with reprogramming
Physical Model	Complex Ginsburg-Landau Equation	Kuramoto Equation
Intrinsic Noise	Yes	No
Controllable Variation	Yes	Yes
Maximum Nodes (projected)	~1000 (with NEMS+ASIC, separated chips)	~1000 (with ~10^6 logic elements Altera Stratix)

III. NEMS-VLSI System-on-Chip (SOC) Architectures

Large scale Oscillator Arrays for Network Control Experiments

Under the MURI NetControl program, we are developing NEMS oscillator arrays as a flexible building block for assembling complex networks. Given their high frequency operation, piezoelectric NEMS oscillators provide fast response to changes in network coupling while providing unprecedented control of nonlinearity, frequency, and the phase and amplitude of coupling. Our current efforts employ printed circuit board (PCB) architectures to connect from 6 to up to 32 oscillators. To scale up the network size beyond 100 oscillators, and possibly up to many thousands, we have proposed to realize all components of the PCB architecture using the System on Chip (SOC) approach. Elements to be integrated include the NEMS resonator arrays, oscillator feedback and signal processing and network control and data processing. Below we describe our cost-effective, three-chip approach to achieving this; in addition to the NEMS resonator array chip it also includes an ASIC (application-specific integrated circuit) for oscillator feedback and signal processing, and a FPGA (field-programmable gate array) chip for network control and data processing.

The NEMS-VLSI SOC Platform

The NEMS-VLSI SOC is a powerful and unprecedented paradigm for net-control experiments. There are many possible implementations; we have chosen one that we believe will be the least expensive, the quickest and most efficient to implement, and readily configurable to a host of possible applications in network science. As shown in Fig. 3, our implementation of the NEMS-VLSI SOC will consists of three chips that, respectively, comprise the NEMS array: the NEMS-VLSI array chip, the "Net-ASIC" chip, and a FPGA controller chip. The NEMS chip comprises an array of, in principle, many thousands of NEMS resonators – each with their own electrical contact pads to provide actuation, transduction and control connections of the electromechanical elements. The NEMS chip itself in this embodiment is passive – it does not include active electronic components (i.e., transistors), nor does it provide connections between the individual NEMS
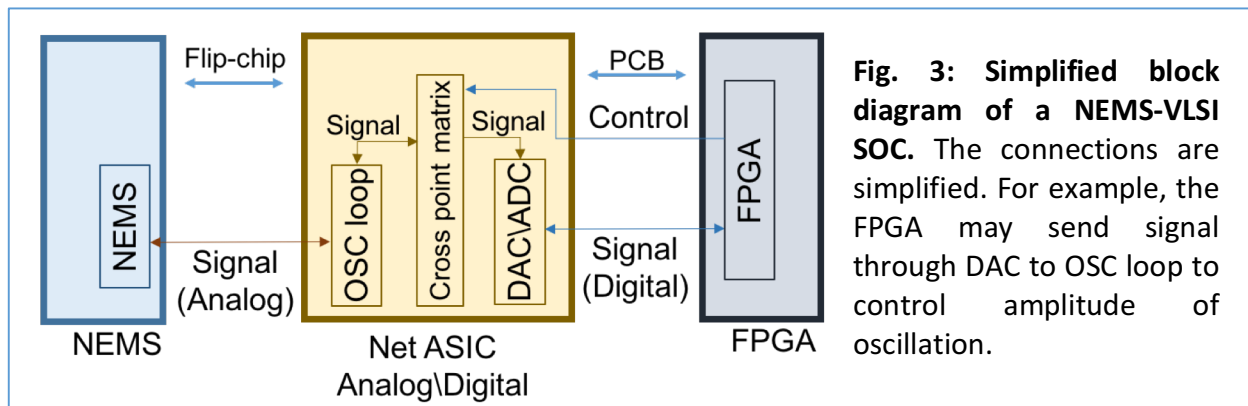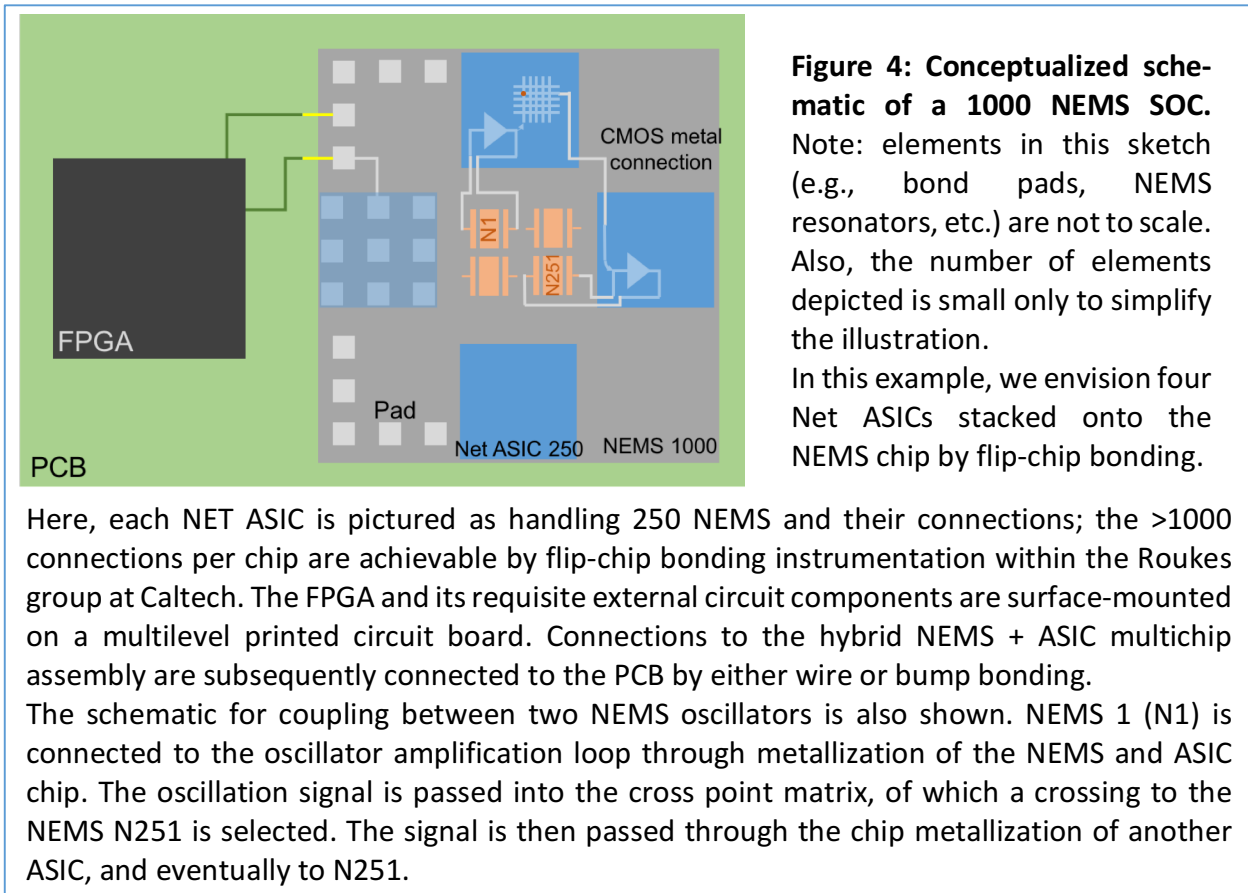

Fig. 3: Simplified block diagram of a NEMS-VLSI SOC. The connections are simplified. For example, the FPGA may send signal through DAC to OSC loop to control amplitude of oscillation.

Figure 4: Conceptualized schematic of a 1000 NEMS SOC.
Note: elements in this sketch (e.g., bond pads, NEMS resonators, etc.) are not to scale. Also, the number of elements depicted is small only to simplify the illustration.
In this example, we envision four Net ASICs stacked onto the NEMS chip by flip-chip bonding.

Here, each NET ASIC is pictured as handling 250 NEMS and their connections; the >1000 connections per chip are achievable by flip-chip bonding instrumentation within the Roukes group at Caltech. The FPGA and its requisite external circuit components are surface-mounted on a multilevel printed circuit board. Connections to the hybrid NEMS + ASIC multichip assembly are subsequently connected to the PCB by either wire or bump bonding.

The schematic for coupling between two NEMS oscillators is also shown. NEMS 1 (N1) is connected to the oscillator amplification loop through metallization of the NEMS and ASIC chip. The oscillation signal is passed into the cross point matrix, of which a crossing to the NEMS N251 is selected. The signal is then passed through the chip metallization of another ASIC, and eventually to N251.

elements. The Net-ASIC chip provides these elements; it has three major components. First, it provides oscillator amplification loops for each of the resonators on the NEMS array chip; these sustain oscillation, maintain a fixed, settable oscillation amplitude, and permit tuning of frequency. Second, it provides a cross-point matrix the permits configuring the oscillator network in real time. Third, it provides both DAC (digital-to-analog conversion) and ADC (analog-to-digital conversion) unit for conversion of waveforms from the oscillator array elements, and for communicating control signals to them in digital and analog formats, as appropriate. The Net ASIC is subsequently interfaced to a controller circuit based on a commercially-available FPGA chip, which will provide both the hub for digital data retrieval, as well as high speed control and modification of the NEMS array network. The FPGA also acts as general interface to the computer using the USB 3.0 standard.

Importantly, all technologies required for creating the NEMS SOC networks are presently both available and well-validated. The NEMS chips will be designed at Caltech by the Roukes group following procedures developed over two decades at Caltech – and subsequently partially fabricated at CEA/LETI in Grenoble, then completed by the Roukes group at Caltech's Kavli Nanoscience Institute. NEMS-VLSI is an advanced, well-validated platform for these experiments. Large scale passive NEMS arrays of up to 1 million elements have been demonstrated by the Alliance for Nanosystems VLSI, a partnership between Caltech and CEA/LETI in Grenoble, France, co-founded by Prof. Roukes in 2006. In late 2007, the Alliance demonstrated the world's first

"standard NEMS process" at the 200mm (8") wafer-scale, which yielded device (NEMS resonator) densities of approximately 1.2×106 cm-2. This has since enabled NEMS devices for applications from high-performance microscale gas analyzers (now commercialized), chip-based mass spectrometers, and ultrahigh frequency signal processing components.

The Net ASIC, although it is custom-designed for the specific NEMS SOC architectures to be pursued, will be manufactured using standard CMOS processes in commercial foundries. It will be designed in collaboration with Prof. Ken Shepard of Columbia University and his research group, who are experts in VLSI design. It will then be manufactured by either GlobalFoundries (née, IBM) facility in New York, or by TSMC in Taiwan.

As previously mentioned, the requisite high performance FPGA controller chips are available commercially.

A 1000-element NEMS SOC

A conceptual schematic of a 1000-element NEMS oscillator array is depicted in Fig. 4. It illustrates the principles we will implement to forming the oscillator networks. The amplification loop of the ASIC and the NEMS form the individual oscillators ("nodes") of the network. Each oscillator signal is routed through a cross point matrix and is subsequently connected to selected oscillators within the network to configure its interconnections ("edges").

A NEMS chip with 1000 or more piezoelectric aluminum nitride (AlN) NEMS elements will be made by pre-processing at CEA/LETI, followed by post-processing, at Caltech, using custom-designed precursor wafers fabricated at CEA/LETI. The CMOS Net ASIC chips, which will be designed at Columbia (in collaboration with Caltech) and manufactured at a commercial foundry, will be flip-chip bonded to the NEMS. The NEMS chip will not only have the NEMS resonator array elements, but (somewhat analogous to a "mother board") will also have connection circuitry to interface the NEMS array to the Net-ASIC chips. The number of bonds per ASIC chip is >1000; this level of complexity is easily managed with technology and expertise already within the Roukes group at Caltech. This NEMS/ASIC chip assembly will then be anchored onto a custom, multi-level PCB along with the FPGA and its requisite ancillary components. The NEMS chip assembly will be connected to this PCB, either by wire bonding or by bump-bonding to the PCB; in this case the requisite number of bonds to between the chip assembly and the PCB is of order several hundred. A local vacuum will be applied to the NEMS array portion at the center of the NEMS-VLSI chip, exactly as has been implemented with NEMS resonators in our already-implemented PCB array architecture. The SOC PCB will be attached to active cooling and temperature control elements to ensure optimal frequency stability of the NEMS oscillators.

Prospective Timeline for Realizing Complex NEMS-VLSI SOCs

The production of the first generation of NEMS-VLSI SOC, which will comprise up to 1000 NEMS oscillators, is expected to take ~18 months. Tasks involved include fabrication and testing of

smaller-scale precursors of the full architecture; design, tape-out and production of chips, programming of FPGA controller; and initial characterization of the complete system.

A subsequent generation of of the NEMS-VLSI SOC architecture that could, in principle, permit upscaling the number of oscillator up to 10,000 nodes is feasible, and will be pursued depending on the result of the first generation, progress of experiments, continued availability of resources, and feedback from initial results.